SCIENCE AND ITS LIMITS: The Regulator's Dilemma

Alvin M. Weinberg

PROLOGUE: *The shift in environmental concerns from visible pollution to more subtle threats, such as toxic pollutants, presents special problems for regulators who must function outside the limits of scientific certainty. The same handicap besets judges who must adjudicate disputes over claims for damages arising from new and hazardous technologies that involve adverse health effects that are latent or unpredictable.*

In this area of uncertainty in which accidental exposure to hazards is rare, scientists resort to probabilistic risk assessment to estimate the likelihood and consequences of events that may carry a threat to human health. Such scientific techniques for the investigation of rare events, however, often cannot provide definitive answers for regulators and judges.

In this essay physicist Alvin Weinberg suggests that instead of asking scientists for answers to unanswerable questions, regulators should settle for less-definitive answers and regulate on the basis of uncertainty. Technological fixes, including greater reliance on inherent safety features that depend on the immutable laws of nature, can help reduce risk. But ultimately, says Weinberg, it may be necessary to establish some threshold beyond which blame for accidents and other untoward events would be unprovable and victims would be compensated by a society as a whole.[1]

Alvin M. Weinberg received his Ph.D. in physics from the University of Chicago in 1939. He has been a leading figure in the development of nuclear energy and has served as director of the Oak Ridge National Laboratory and as director of the Institute for Energy Analysis of the Oak Ridge Associated Universities. He is the coauthor of The Physical Theory of Neutron Chain Reaction (1958) and has written extensively on nuclear energy, nuclear proliferation, and the interaction between modern technology and society.

In his essay "Risk, Science, and Democracy," William D. Ruckelshaus expresses very clearly what I call the regulator's dilemma. During the past 15 years, Ruckelshaus notes, there has been a shift in public emphasis from visible and demonstrable pollution problems, such as smog resulting from automobiles and raw sewage, to potential and largely invisible problems, such as the effects of low concentrations of toxic pollutants on human health. This shift is important for two reasons. First, it has changed the way that science is applied to practical questions of public health protection and environmental regulation. Second, it has raised difficult questions about managing chronic risks within the context of free and democratic institutions.[2]

When the environmental concern was patent and obvious—such as the problem of smog in Los Angeles—science could and did provide unequivocal answers. Smog, for example, comes from the gas emissions from burning liquid hydrocarbons, and the answer to the smog problem lies in controlling these emissions. The regulator's course was rather straightforward because the science upon which regulatory decisions are made was operating well within its power. However, when the environmental concern is subtle—for example, how much cancer is caused by an increase of 10 percent in mean background radiation—science is being asked a question that lies beyond its power; the question is trans-scientific. Yet the regulator, by law, is expected to regulate even though science can hardly help him; this is the regulator's dilemma.

Although my essay is subtitled The Regulator's Dilemma, many of the same issues arise in the adjudication of disputes over who is to blame and who is to be compensated for damage allegedly caused by rare events, such as nuclear accidents. The regulator's dilemma is also faced by the judge who is presiding over a tort case involving, for example, a claim for damages blamed on a toxic waste dump. Indeed, the regulator's dilemma could equally be called the toxic tort dilemma.

A lawsuit involving alleged injury from chemical pollutants is unlike the traditional liability case. If my car injures a pedestrian, I am liable to be sued. What is at issue, however, is not whether I have injured a pedestrian. Rather, it is whether I am at fault. On the other hand, if the lead from my car's exhaust is alleged to cause bodily harm, the issue is not whether my car emitted the lead but whether the lead actually caused the alleged harm. The two situations are quite different. In the first example the relation between cause and injury is not at issue. In the second it is the issue.

In this essay, therefore, I try to delineate more precisely those limits to science that give rise to the regulator's dilemma. I speculate on how these intrinsic limits to science seem to have catalyzed a profound attack on science by some sociologists and public-interest activists. In addition, I offer a few ideas that may help the harried regulators finesse these trans-scientific issues.

II

Science deals with regularities in our experience; art deals with singularities. It is no wonder that science tends to lose its predictive or even explanatory power when the phenomena it deals with are singular,

2025545699

irreproducible, and one of a kind—in other words, rare. Although science can often analyze a rare event after the fact—for example, the extinction of dinosaurs during the Cretaceous-Tertiary period following the presumed collision of the earth and an asteroid—it has great difficulty predicting when such an uncommon event will occur.

I distinguish here between two sorts of rare events—accidents and low-level insults, whose potential to cause injury is unknown. Accidents are large-scale malfunctions whose etiology is not in doubt, but whose likelihood is very small. The partial nuclear reactor meltdown at Three Mile Island in 1979 and the release of toxic gas from a chemical plant at Bhopal, India, in 1984 are examples of accidents. The precursors to these specific events—for example, the condition of the auxiliary water feed system and other components at Three Mile Island—and the way in which the accidents unfolded are well understood. Estimates of the likelihood of the particular sequence of malfunctions are less firmly grounded. As the number of individual accidents increases, prediction of their probability becomes more and more reliable. We can predict very well how many automobile fatalities will occur in 1986; we can hardly claim the same degree of reliability in predicting the number of serious reactor accidents in 1986.

Low-level insults are rare in a rather different sense. We know that about 100 rems of radiation will double the mutation rate in a large population of exposed mice. How many mutations will occur in a population of mice exposed to 100 millirems of radiation? In this case the mutations, if induced at all by such low levels of exposure, are so rare that to demonstrate an effect with 95 percent confidence would require the examination of many millions of mice. Although such an effort is not impossible in principle, it is in practice. Moreover, even if we could perform so heroic a mouse experiment, the extrapolation of these findings to humans would still be fraught with uncertainty. Thus, human injury or abuse from low-level exposure to radiation is a rare event whose frequency cannot be accurately predicted.

III

When dealing with events of this sort, science resorts to the language of probability. Instead of saying that this accident will happen on that date, or that a particular person exposed to a low-level dose of radiation will suffer a particular fate, it tries to assign probabilities for such occurrences. Of course, where the number of instances is very large or the underlying mechanisms are fully understood, the probabilities are themselves perfectly reliable. In quantum mechanics there is no uncertainty as to the probability distribution of the phenomenon being described. In the class of phenomena considered here, however, even though the likelihood of an event happening or of a disease being caused by a specific exposure is given as a probability, the probability itself is very uncertain. One can think of a somewhat fuzzy demarcation between what I have called science and trans-science. The domain of science covers phenomena that are deterministic or whose probability of occurrence can itself be stated precisely; in contrast, trans-science covers those events whose probability of occurrence is itself highly uncertain.

Although science can often analyze a rare event after the fact, it has great difficulty predicting when such an uncommon event will occur.

Despite the difficulties, scientific mechanisms have been devised for estimating, however imperfectly, the probability of rare events. For accidents the technique is probabilistic risk assessment (PRA); for low-level insults various empirical and theoretical approaches are used.

Although probabilistic risk assessment had been used in the aerospace industry for a long time (for example, to predict the reliability of components), it first sprang into public prominence in 1975 with a reactor safety study directed by nuclear engineer Norman C. Rasmussen.[3] The Rasmussen study, sponsored by the Atomic Energy Commission (now known as the Nuclear Regulatory Commission), was designed to estimate the public risks involved in potential accidents at commercial nuclear reactors.

Probabilistic risk assessment, when applied to nuclear reactors, seeks to identify all sequences of subsystem failures that may lead to a failure of the overall system; it then tries to estimate the consequences of each subsystem failure so identified. The result is a probability distribution, $P(C)$; that is, the probability, P, per reactor year, of a consequence having magnitude C. Consequences include both material damage and health effects. Usually, the probability of accidents having large consequences is less than the probability of accidents having small consequences.

A probabilistic risk assessment for a reactor requires two separate estimates: first, an estimate of the probability of each accident sequence; second, an estimate of the consequences—particularly the damage to human health—caused by the uncontrolled radioactive effluents released in the accident. An accident sequence is a series of equipment or human malfunctions, such as a pump that fails to start, a valve that does not close, or an operator confusing an ON with an OFF signal. We have statistical data for many of these individual events; for example, enough valves have operated for enough years so that we can, at least in principle, make pretty good estimates of the probability of failure.

Uncertainties still remain, however, because we can never be certain that we have identified every relevant sequence. Proof of the adequacy of probabilistic risk assessment must therefore await the accumulation of operating experience. For example, the median probability of a core melt in a light water reactor, according to the 1975 Rasmussen study, was 1 in every 20,000 reactor-years; the core melt at Three Mile Island's number two reactor (TMI-2) occurred after only 700 reactor-years. The number two reactor, however, differed from the reactors Rasmussen studied, and in retrospect, one could rationalize most of the discrepancy between his estimate and the seemingly premature occurrence at TMI-2.

Since the core melt at Three Mile Island, the world's light water reactors have accumulated some 1,500 reactor-years of operation without a core melt. This performance places an upper limit on the a priori estimate of the core-melt probability. Thus, if this probability were as high as 1 in every 1,000 reactor years, the likelihood of surviving 1,500 reactor-years would not be more than 22 percent; put otherwise, we can say with 78 percent confidence that the core-melt probability is not as high as 1 in 1,000 reactor years. With 500 light water reactors on line in the world, should we survive until the year 2000 without another core melt, we could then say with 95 percent confidence

ISSUES IN SCIENCE AND TECHNOLOGY

that the core-melt probability is not higher than 1 in 3,000 reactor-years. In the absence of such experience, one is left with rather subjective judgments. Although Harold W. Lewis, in his critique of Rasmussen's 1975 study,[4] asserts that he could not place a bound on the uncertainty of probabilistic risk assessment, Rasmussen argued that his estimate of core-melt probability may be in error by about a factor of 10 either way—that is, the probability may be as high as 1 in 2,000 reactor-years or as low as 1 in 200,000 reactor-years.

As we see, after 3,000 reactor-years of operation without a core melt, we can say with about 78 percent confidence that Rasmussen's upper limit (1 in 2,000 reactor-years) is not too optimistic. Furthermore, if we survive to the year 2000 without a core melt, the confidence level with which we can make this assertion rises to 95 percent. Our confidence in probabilistic risk assessment can eventually be tested against actual, observable experience. Until this experience has been accumulated, however, we must concede that any probability we predict must be highly uncertain. To this degree our science is incapable of dealing with rare accidents, but time, so to speak, annihilates uncertainty in estimates of accident probability.

Unfortunately, time does not annihilate uncertainties over consequences as unequivocally as it does uncertainties over frequency of accidents. A large reactor or chemical plant accident can cause both immediate, acute health effects and delayed, chronic effects. If the exposure either to radiation or to methyl isocyanate is high enough, the effect on health is quite certain. For example, a single exposure of about 400 rems will cause about half of the people exposed to die. On the other hand, in a large accident many people will also be exposed to smaller doses—indeed, to doses so low that the resulting health effects are undetectable. At Bhopal many thousands of people were exposed to methyl isocyanate but they recovered. We cannot say positively whether or not they will suffer some chronic disability.

The very worst accident envisaged in the Rasmussen study, with a probability of 1 in 1 billion reactor-years, projected an estimated 3,300 early fatalities, 45,000 early illnesses, and 1,500 delayed cancers per year among 10 million exposed people. Almost all of the estimated delayed cancers are attributed to exposures of less than 1,000 millirems per year—a level at which we are very hard put to estimate the risk of inducing cancer. Similarly, the American Physical Society's critique of the Rasmussen study attributed an additional 10,000 deaths over 30 years among 10 million people exposed to cesium-135 distributed in a very large accident.[5] The average exposure in this case was assumed to be 250 millirems per year—again, a level at which our estimates of the health effects are extremely uncertain.

Has the nuclear community, particularly its regulators, figuratively shot itself in the foot by trying to estimate the number of delayed casualties as a result of these low-level exposures? In retrospect, I think the Rasmussen study would have been on more solid ground had it confined its estimates to those health effects resulting from exposures at higher levels, where science makes reliable estimates. For the lower exposures the consequences could have been stated simply as the number of man-rems (the number of people multiplied by the number of rems) of exposure of individuals whose total exposure did not exceed, say, 5,000 millirems, without trying to convert this man-rems number

Our confidence in probabilistic risk assessment can eventually be tested against actual, observable experience.

2025545702

into numbers of latent cancers. Thus, health consequence would be reported in two categories: (1) for highly exposed individuals, the number of health effects; and (2) for slightly exposed individuals, the total man-rems or even the distribution of exposures accrued by the large number of individuals so exposed. Perhaps such a scheme could be adopted in reporting the results of future probabilistic risk assessments; at least it has the virtue of being more faithful than the present convention to the state of scientific knowledge

<div align="center">IV</div>

In both of my examples of accidents (Bhopal and nuclear accidents), many people are exposed to low-level insult. The uncertainties inherent in estimating the effects of such low-level exposure are heaped on top of the uncertainties in estimating the probability of the accident that may lead to exposure in the first place.

Science has exerted great effort to ascertain the shape of the dose-response curve at low dose—but very little, if anything, can be said with certainty about the low-dose response. Thus, to quote the report of the National Research Council, *The Effects on Populations of Exposure to Low Levels of Ionizing Radiation: 1980* (also known as BEIR-III, for the committee that prepared it, the Committee on the Biological Effects of Ionizing Radiation), "The Committee does not know whether dose rates of gamma or x-rays of about 100 mrads/yr are detrimental to man. . . . It is unlikely that carcinogenic and teratogenic effects of doses of low-LET radiation administered at this dose rate will be demonstrable in the foreseeable future."[6] This prompted Philip Handler, then president of the National Academy of Sciences, to comment in his letter of transmittal to the Environmental Protection Agency, which had requested the study, "It is not unusual for scientists to disagree . . . (and) . . . the sparser and less reliable the data base, the more opportunity for disagreement. . . . The report has been delayed . . . to permit time . . . to display all of the valid opinions rather than distribute a report that might create the false impression of a clear consensus where none exists."[7]

This forthright admission that science can say little about low-level insults I find admirable. It represents an improvement over the unjustified assertion in the BEIR-II report of 1972 that 170 millirems per year over 30 years, if imposed on the entire U.S. population, would cause between 3,000 and 15,000 cancer deaths per year.[8] I do not quarrel with the estimated upper limit—which amounts to 1 cancer per 2,500 man-rems, but I regard placing the lower limit at 3,000 rather than at zero as unjustified. Moreover, I think it has caused great harm. The proper statement should have been that at 170 millirems per year, we estimate the upper limit for the number of cancers to be 15,000 per year; the lower limit may be zero.

Since the appearance of the BEIR reports, two other developments have added to the burden of those who must judge the carcinogenic hazard of low-level insults: an awareness and study of (1) natural carcinogens, and (2) ambiguous carcinogens.

Natural carcinogens. Is cancer environmental in the sense of being caused by technology's effluents, or is it a natural consequence of aging? In the

past few years we have seen a remarkable shift in viewpoint; whereas 15 years ago most cancer experts would have accepted a primarily environmental etiology for cancer, today the view that natural carcinogens are far more important than are manmade ones has gained many converts. In his 1983 article in *Science*, biochemist Bruce N. Ames marshaled powerful evidence that many of our most common foods contain naturally occurring carcinogens.[9] Indeed, biochemist John R. Totter, former director of the Atomic Energy Commission's division of biology and medicine, has offered evidence for the oxygen radical theory of carcinogenesis: that we eventually get cancer because we metabolize oxygen and subsequently produce oxygen radicals that can play havoc with our DNA.[10] As such views of the etiology of cancer acquire scientific support, I think that the trans-scientific question, as to how much cancer is caused by a tiny chemical or physical insult will be recognized as irrelevant. One does not swat gnats when pursued by elephants.

Ambiguous carcinogens. To further complicate the cancer picture, there is evidence that some agents, such as dioxin, various dyes, and even moderate levels of radiation, seem to diminish the incidence of some cancers while simultaneously increasing the incidence of others. The lifespan of the animals exposed to these agents in laboratory tests on average exceeds that of animals not exposed.[11] A most striking example, given by biostatistician Joseph K. Haseman, is yellow dye number 14 given to leukemia-prone female rats. This dye completely suppresses leukemia, which is always fatal, but causes liver tumors, most of which are benign.

I mention these two findings—or perhaps they should be considered points of view—to stress my underlying point: that when we are concerned with low-level insult to human beings, we can say very little about the cancer dose-response curve. Saying that so many cancers will be caused by so much low-level exposure to so many people, a practice that terrifies many people, goes far beyond what science actually can say.

Some agents, such as dioxin, various dyes, and even moderate levels of radiation, seem to diminish the incidence of some cancers while simultaneously increasing the incidence of others.

V

Does the scientific community accept the notion that there are intrinsic limits to what it can say about rare events; that as events become rarer, the uncertainty in the probability of occurrence of a rare event is bound to grow? Perhaps a better way of framing this question is: Of what use can we put scientific tools of investigation of rare events, such as probabilistic risk assessment and large-scale animal experiments, if we concede that we can never get definitive answers?

I believe that probabilistic risk assessment with an uncertainty factor as high as 10 is often useful, especially if one uses the technique for comparing risks. For example, the 1,500 reactor-years already experienced since the Three Mile Island accident suggest that a reactor core-melt probability is likely to be less than 1 in 1,000 reactor-years and may well be as low as less than 1 in 10,000 reactor-years. This is to be compared with dam failures whose probability, based on many hundreds of thousands of dam-years (where time has annihilated uncertainty), is around 1 in 10,000 dam-years. Even with an uncertainty factor of 10, we can judge how safe reactors are compared to dams.

When one compares the relative intrinsic safety of two very similar devices—such as two water-moderated reactors—probabilistic risk assessment is on much more solid ground. Here one is not asking for absolute estimates of risk, but rather for estimates of relative safety. If reactors A and B differ in only a few details—say reactor A has two auxiliary water feed trains whereas B has only one—the ratio of core-melt probabilities should be much more reliable than their absolute values because the ratio requires an estimate of failure of a single subsystem, in this instance the extra auxiliary water feed on reactor A.

Not only can one say with reasonable assurance how much safer reactor A is than reactor B, but as a result of the detailed analysis one can identify the subsystems that contribute most to the estimated failure rate. Even if probabilistic risk assessment is inaccurate, it is very useful in unearthing deficiencies; one can hardly deny that a reactor in which deficiencies revealed by probabilistic risk assessment have been corrected is safer than one in which they have not been corrected, even if one is unwilling to say how much safer.

The Delaney clause is the worst example of how a disregard of an intrinsic limit of science can lead to bad policy by overenthusiastic politicians.

Somewhat the same considerations apply to low-level insult. An agent that does not shorten lifespan at high dose will not shorten lifespan at low dose. An agent that is a very powerful carcinogen at high dose is more likely to be a carcinogen at low dose than one that is a less powerful high-dose carcinogen. Thus, animal experiments surely are useful in deciding which agents to worry about and which not to worry about. Of course, the Ames test (which determines by a relatively simple procedure whether a substance is mutagenic) has made at least some preliminary screening of carcinogens more feasible because substances that cause mutations are considered to be potential carcinogens. The difficulty today seems to be not so much identifying agents that at high dose may be carcinogens as it is prohibiting exposures far below levels at which no effect can be, or perhaps ever will be, demonstrated. The regulator and the concerned citizen are inclined to approve the Delaney clause of the Federal Food, Drug, and Cosmetic Act, which prohibits the use of any food additive that has been shown to cause cancer in laboratory animals or humans. This clause, however, is of no help in resolving such issues as the relative risks of, say, cancer induction by nitrosamines (carcinogenic compounds that can be formed in the body from nitrites) and digestive disorders caused by meat untreated with nitrites.

The Delaney clause is the worst example of how a disregard of an intrinsic limit of science can lead to bad policy by overenthusiastic politicians. Harvard physicist Harvey Brooks has often pointed out that one can never prove the impossibility of an event that is not forbidden by a law of nature. Most will agree that a perpetual motion machine is impossible because it violates the laws of thermodynamics. That one molecule of a polychlorinated biphenyl (PCB) may cause a cancer in humans is a proposition that violates no law of nature: hence many, even within the scientific community, seem willing to believe that this possibility is something to worry about. It was this error that led to the Delaney clause.

When is an event so rare that the prediction of its occurrence forever lies outside the domain of science and therefore within the domain of trans-science? Clearly we cannot say, and perhaps as science progresses, this

ISSUES IN SCIENCE AND TECHNOLOGY

2025545705

boundary between science and trans-science will recede toward events of lower frequency. At any stage, however, the boundary is fuzzy, and much scientific controversy boils over deciding where it lies. One need only read the violent exchange between Edward P. Radford and Harald H. Rossi over the risk of cancer from low levels of radiation to recognize that where the facts are obscure, argument—even ad hominem argument—blossoms.[12] Indeed, Alice Whittemore in her "Facts and Values in Risk Analysis for Environmental Toxicants," has pointed out that facts and values are always intermingled at this "rare event" boundary between science and trans-science.[13] A scientist who believes that nuclear energy is evil because it inevitably leads to proliferation of nuclear weapons (which is a common basis for opposition to nuclear energy) is likely to judge the data on induction of leukemia from low-level exposures at Nagasaki differently than is a scientist whose whole career has been devoted to making nuclear power work. Cognitive dissonance is all but unavoidable when the data are ambiguous and the social and political stakes are high.

VI

No one would dispute that judgments of scientific truth are much affected by the scientist's value system when the issues are at or close to the boundary between science and trans-science. On the other hand, as the matter under dispute approaches the domain of science, most would claim that the scientist's extrascientific values intrude less and less. Soviet scientists and U.S. scientists may disagree on the effectiveness of a ballistic missile defense, but they agree on the cross section of U^{235} or the lifetime of the pi meson.

This all seems obvious, even trite. Yet in the past decade or so a school of sociology of knowledge has sprung up in Great Britain that claims that "scientific views are determined by social (external) conditions, rather than by the internal logic of scientific tradition and inherent characteristics of the phenomenal world,"[14] or that "all knowledge and knowledge claims are to be treated as being socially constructed: genesis, acceptance, and rejection of knowledge [is] sought in the domain of the Social World rather than . . . the Natural World."[15]

The attack here is not on science at the boundary with trans-science, in particular—the prediction of the frequency of rare events. At least the more extreme of the sociologists of knowledge claim that using traditional ways of establishing scientific truth—by appealing to nature in a disciplined manner—is not how science really works. Scientists are seen as competitors for prestige, pay, and power, and it is the interplay among these conflicting aspirations, not the working of some underlying scientific ethic, that defines scientific truth. To be sure, these attitudes toward science are not widely held by practicing scientists; however, they are taken seriously by many political activists who, though not in the mainstream of science, nevertheless exert important influence on other institutions—the press, the media, the courts—that ultimately influence public attitudes toward science and its technologies.

If one takes such a caricature of science seriously, how can one trust a scientific expert? If scientific truth, even at the core of science, is decided by

negotiation between individuals in conflict because they hold different non-scientific beliefs, how can one say that this scientist's opinion is preferable to that one's? Furthermore, if the matter at issue moves across the boundary between science and trans-science, where all we can say with certainty is that uncertainties are very large, how much less able are we to distinguish between the expert and the charlatan, between the scientist who tries to adhere to the usual norms of scientific behavior and the scientist who suppresses facts that conflict with his political, social, or moral preconceptions?

One way to deal with these assaults on scientists and scientific truth would be to define a new branch of science, called regulatory science, in which the norms of scientific proof are less demanding than are the norms in ordinary science. I should think that a far more honest and straightforward way of dealing with the intrinsic inability of science to predict the occurrence of rare events is to concede this limitation and not to ask of science or scientists more than they are capable of providing. Instead of asking science for answers to unanswerable questions, regulators should be content with less far-reaching answers. For example, where the ranges of uncertainty can be established, regulate on the basis of uncertainty; where the ranges of uncertainty are so wide as to be meaningless, recast the question so that regulation does not depend on answers to the unanswerable. Furthermore, because these same limits apply to litigation, the legal system should recognize, much more explicitly than it has, that science and scientists often have little to say, probably much less than some scientific activists would admit.

The expertise of scientific adversaries is often at the heart of litigation over personal injury alleged to be caused by subtle, low-level exposures. Each side presents witnesses whose scientific credentials it regards as impeccable. Because the issues themselves tend to be trans-scientific, one can hardly decide the validity of the assertions of either side's witnesses. Under the circumstances, I suppose, one is justified in regarding a scientific witness no differently than any other witness; his credibility is judged by his past record, behavior, and general demeanor, as well as the self-consistency of his testimony. Such, at least, was the way in which a federal district court judge, Patrick Kelley, settled *Johnston v. United States*, in which the issue was the claim that exposure to radiation from reworking old aircraft instrument dials had caused injury; Kelley impugned, on grounds no different from those one would invoke in an ordinary lawsuit, the competence if not the integrity of some of the plaintiff's scientific witnesses.

VII

There are various ways to provide some assurance of safety despite uncertainty. Here I briefly describe two of these ways—which I call the technological fix and de minimis—without claiming that these are the most important, let alone the only, ones.

Technological fix. Science cannot exactly predict the probability of a serious accident in a light water reactor or the likelihood that a radioactive waste canister in a depository will dissolve and release radioactivity to the environment. Can one design reactors or waste canisters for which the

Instead of asking science for answers to unanswerable questions, regulators should be content with less far-reaching answers.

probability of such occurrences is zero—or at least, where the prevention of such mishaps relies on immutable laws of nature that can never fail rather than on the less than reliable intervention of electromechanical devices? Surprisingly, this approach to nuclear safety has come into prominence only in the past five years. Kåre Hannerz in Sweden and Herbert Reutler and Günter H. Lohnert in West Germany have proposed reactor systems whose safety does not depend on active interventions, but rather on passive, inherent characteristics.[16] Although one cannot say that the probability of mischance has been reduced to zero, there is little doubt that the probabilities are several, perhaps three, orders of magnitude lower than the probabilities of mischance for existing reactors. To the extent that such proposed reactors embody the principle of inherent safety, their adoption would avoid much of the dispute over reactor safety, the limits on nuclear accident liability contained in the Price-Anderson Act, repetition of the Three Mile Island accident, and so forth. In short, such a technological fix enables one largely to ignore the uncertainties in any prediction of core-melt probabilities.

The idea of incorporating inherent or passive safety into the design of chemical plants had been proposed by Trevor A. Kletz of the Loughborough University of Technology in 1974, shortly after the disaster at the Flixborough cyclohexane plant, which killed 28 people.[17] I suspect that one of the main consequences of the Bhopal disaster will be the incorporation of inherent safety features into new chemical plants; again, a way of finessing uncertainty in predicting failure probabilities.

De minimis. A perfect technological fix, such as a totally safe reactor or a crash-proof car, is usually not available, at least at an affordable cost. Some low-level exposure to materials that are toxic at high levels is inevitable, even though we can never accurately establish the risk of such exposure. One way of dealing with this situation is to invoke the principle of de minimis. This principle, as Howard I. Adler and I suggested several years ago, argues that for insults that occur naturally and to which the biosphere has always been exposed and presumably to which it has adapted, one should not worry about any additional man-made exposure as long as the man-made exposure is small compared to the natural exposure.[18] The basic idea is that the natural level of a ubiquitous exposure (such as cosmic radiation), if it is deleterious, cannot have been very deleterious because in spite of its ubiquity, humans have survived. Moreover, we do not know and can never know what the residual effect of that natural exposure really is. An additional exposure that is small compared to natural background radiation should be acceptable; at the very least, its deleterious effect, if any, cannot be determined.

Adler and I suggested that for radiation whose natural background is well known, one may choose a de minimis level as the standard deviation of the natural background. This turns out to be around 20 percent of the mean background, around 20 millirems per year; this value has been used as the Environmental Protection Agency standard for exposure to the entire radio-chemical fuel cycle.

Scientists know more about the natural incidence and biological effects of radiation than they do about any other agent. It would be natural, therefore, to use the standard established for radiation as a standard for other agents.

This approach has been used by chemist T. Westermark of the Royal Institute of Technology in Stockholm. He has suggested that for naturally occurring carcinogens such as arsenic, chromium, and beryllium, one may choose a de minimis level to be, say, 10 percent of the natural background.[19]

Clearly, a de minimis level will always be somewhat arbitrary. Nevertheless, it seems to me that unless such a level is established, we shall forever be involved in fruitless arguments, the only beneficiaries of which will be the toxic tort lawyers. Could the principle of de minimis be applied in litigation in much the same way it may be applied to regulation—that is, if the exposure is below de minimis, then the blame is intrinsically unprovable and cannot be litigated? I would imagine that the legal de minimis may be set higher than the regulatory de minimis; for example, the legal de minimis for radiation could be the background (after all, the BEIR-III committee concedes there is no way of knowing whether or not such levels are deleterious). The regulatory de minimis could justifiably be lower, simply on grounds of erring on the side of safety.

One approach may be to concede that there is some level of exposure that is beyond demonstrable effect. This defines a trans-scientific threshold. A de minimis level could then be established at some fraction, say one-tenth, of this beyond-demonstrable-effect level. For example, if we take 100 millirems per year of radiation as the beyond-demonstrable-effect level for general somatic effects (damaging somatic cells as opposed to germline cells), which is the value according to the BEIR-III committee, a de minimis level could be set at 10 millirems per year. Of course, such a procedure would evoke much controversy as to what is the beyond-demonstrable-effect level or whether 10 is an ample safety factor. This example demonstrates, however, that at least in the case of low-level radiation, a scientific committee has been able to agree on a beyond-demonstrable-effect level. As for the safety factor of 10, this cannot be adjudicated on scientific grounds. The most one can say is that tradition often supports a safety factor of 10—for example, the old standard for public exposure (500 millirems per year) was set at one-tenth of the tolerance level for workers (5,000 millirems per year).

Can the principle of de minimis be applied to accidents? What I have in mind is the notion that accidents that are sufficiently rare may be regarded somehow in the same category as acts of God and be compensated accordingly. We already recognize that natural disasters should be compensated by the society as a whole. One can argue that an accident whose occurrence requires an exceedingly unlikely sequence of untoward events may also be regarded as an act of God. Thus, the Price-Anderson Act could be modified so that, quite explicitly, accidents whose consequences exceeded a certain level, and whose probability as estimated by probabilistic risk assessment would be less than, say, 1 in 1 billion per year, would be treated as acts of God. Compensation in excess of the amount stipulated in the revised act would be the responsibility of Congress. The cutoff for either compensation or for probabilities would be negotiable, and perhaps it would be revised every 10 years or so. One not entirely fanciful suggestion may be to set any probability of the order of 1 in 10 million to 1 in 100 million per year to be a de minimis cutoff, this being the frequency at which the earth may have been visited by

One can argue that an accident whose occurrence requires an exceedingly unlikely sequence of untoward events may also be regarded as an act of God.

ISSUES IN SCIENCE AND TECHNOLOGY

2025545709

the cometary asteroids that may have caused the extinction of species in past geologic eras.

VIII

As in most such questions, identifying and characterizing the problem is easier than solving it. That the dilemma of the regulator and the toxic tort judge is rooted in science's inability to predict rare events cannot be denied. Getting the regulator and the toxic tort judge off the horns of the dilemma is far from easy, and my two suggestions—the technological fix and de minimis—are offered tentatively and with diffidence.

Equally obvious is the intrinsic social dimension of the issue. In an open, litigious democracy such as ours, any regulation and any judicial decision can be appealed, and if the courts offer no redress, Congress, in principle, can do so. These legal mechanisms are ponderous, however. The result seems to me to be a gradual slowing of our technological-social engine as we become more and more enmeshed in fruitless argument over unresolvable questions.

Western society was debilitated once before by such fruitless tilting with Don Quixotian windmills. I refer of course to the devastating campaign against witches from the fourteenth century to the early seventeenth century. As ecologist William Clark has put it so vividly, society took it for granted during that period that death, disease, and crop failure could be caused by witches.[20] To avoid such catastrophes, one had to burn the witches responsible for them—and consequently some million innocent people were burned. Finally, in 1610, the Spanish inquisitor Alonzo Salazar y Frías realized there was no demonstrated connection between catastrophe and witches. Although he did not prohibit the burning of witches, he did prohibit use of torture to extract confessions. The burning of witches, and witch hunting generally, declined precipitously.

I have recounted this story many times by now. Yet it still seems to me to capture the essence of our dilemma: the connection between low-level insult and bodily harm is probably as difficult to prove as the connection between witches and failed crops. I regard it as an aberration that our society has allowed this issue to emerge as a serious social concern, which in the modern context is hardly less fatuous than were the witch hunts of the past. That dark phase in western society died out only after several centuries. I hope our open, democratic society can regain its sense of proportion far sooner and can get on with managing the many real problems we always will face rather than waste its energies on essentially insoluble, and by comparison, intrinsically unimportant, problems. ■

NOTES:

1. This article was adapted from a paper delivered at a June 3–4, 1985, National Academy of Engineering symposium on "Hazards: Technology and Fairness." A report on that symposium will be published in book form by the National Academy Press.

2. William D. Ruckelshaus, "Risk, Science, and Democracy," *Issues in Science and Technology* 1 (Spring 1985): 19–38.

3. U.S. Nuclear Regulatory Commission, *Reactor Safety Study: An Assessment of Accident Risk in U.S. Commercial Nuclear Plants* (WASH-1400, NUREG 75/014) (Washington, D.C., 1975).

4. U.S. Nuclear Regulatory Commission, *Risk Assessment Review Group Report to the U.S. Nuclear Regulatory Commission* (NUREG/CR-0400) (Washington, D.C., September 1978), vi.

5. "Report to the American Physical Society by the Study Group on Light Water Reactor Safety," *Reviews of Modern Physics* 47 (Supplement 1) (Summer 1975).

6. National Research Council, *The Effects on Populations of Exposure to Low Levels of Ionizing Radiation: 1980* (BEIR-III), (Washington, D.C.: National Academy Press, 1980), 2.

7. National Research Council, *The Effects on Populations of Exposure to Low Levels of Ionizing Radiation: 1980* (BEIR-III), iii.

8. National Research Council, *The Effects on Populations of Exposure to Low Levels of Ionizing Radiation* (BEIR-II), (Washington, D.C.: National Academy Press, 1972), 2.

9. Bruce N. Ames, "Dietary Carcinogens and Anticarcinogens," *Science* 221 (Sept. 23, 1983): 1249, 1256-64.

10. John R. Totter, "Spontaneous Cancer and Its Possible Relationship to Oxygen Metabolism," *Proceedings of the National Academy of Sciences* 77 (April 1980): 1763-67.

11. Alvin M. Weinberg and John B. Storer, "On 'Ambiguous' Carcinogens and Their Regulation," *Risk Analysis* 5 (June 1985): 151-55.

12. National Research Council, *The Effects on Populations of Exposure to Low Levels of Ionizing Radiation: 1980* (BEIR-III), 287-321.

13. Alice Whittemore, "Facts and Values in Risk Analysis for Environmental Toxicants," *Risk Analysis* 3 (March 1983): 23-33.

14. John Ben-David, "Emergence of National Traditions in the Sociology of Science: The United States and Great Britain," *Sociological Inquiry* 48, nos. 3 and 4 (1978): 209.

15. Trevor J. Pinch and Wiebe E. Bijker, "The Social Construction of Facts and Artefacts: Or How the Sociology of Science and the Sociology of Technology Might Benefit Each Other," *Social Studies of Science* 14 (1984): 401.

16. Kåre Hannerz, *Towards Intrinsically Safe Light Water Reactors* (ORAU/IEA-83-2(M) Rev.) (Oak Ridge, Tenn.: Oak Ridge Associated Universities, Institute for Energy Analysis, June 1983); Herbert Reutler and Günter H. Lohnert, "The Modular High Temperature Reactor," *Nuclear Technology* 62 (July 1983): 22-30.

17. Trevor A. Kletz, *Cheaper, Safer Plants or Wealth and Safety at Work: Notes on Inherently Safer and Simpler Plants* (Rugby, England: The Institution of Chemical Engineers, 1984).

18. Howard I. Adler and Alvin M. Weinberg, "An Approach to Setting Radiation Standards," *Health Physics* 34 (June 1978): 719-20.

19. T. Westermark, *Persistent Genotoxic Wastes: An Attempt at a Risk Assessment* (Stockholm: Royal Institute of Technology, 1980).

20. William C. Clark, *Witches, Floods, and Wonder Drugs: Historical Perspectives on Risk Management* (RR-81-3) (Laxenburg, Austria: International Institute for Applied Systems Analysis, March 1981).